

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants:	Sherif Yacoub, et al.	Examiner:	Leonard Saint Cyr
Serial No.:	10/668,141	Group Art Unit:	2626
Filed:	September 23, 2003	Docket No.:	200300101-1
Title:	System and Method Using Multiple Automated Speech Recognition Engines		

APPEAL BRIEF UNDER 37 C.F.R. § 41.37

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This Appeal Brief is filed in response to the Final Office Action mailed September 17, 2007 and Notice of Appeal filed on December 17, 2007.

AUTHORIZATION TO DEBIT ACCOUNT

It is believed that no extensions of time or fees are required, beyond those that may otherwise be provided for in documents accompanying this paper. However, in the event that additional extensions of time are necessary to allow consideration of this paper, such extensions are hereby petitioned under 37 C.F.R. § 1.136(a), and any fees required (including fees for net addition of claims) are hereby authorized to be charged to Hewlett-Packard Development Company's deposit account no. 08-2025.

I. REAL PARTY IN INTEREST

The real party in interest is Hewlett-Packard Development Company, LP, a limited partnership established under the laws of the State of Texas and having a principal place of business at 20555 S.H. 249 Houston, TX 77070, U.S.A. (hereinafter "HPDC"). HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, LLC.

II. RELATED APPEALS AND INTERFERENCES

There are no known related appeals, judicial proceedings, or interferences known to appellant, the appellant's legal representative, or assignee that will directly affect or be directly affected by or have a bearing on the Appeal Board's decision in the pending appeal.

III. STATUS OF CLAIMS

Claims 1 – 20 are pending in the application and stand finally rejected. The rejection of claims 1 – 20 is appealed.

IV. STATUS OF AMENDMENTS

No amendments were made after receipt of the Final Office Action. All amendments have been entered. Further, no amendments were made to the claims during prosecution. As such, all claims are original.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The following provides a concise explanation of the subject matter defined in each of the claims involved in the appeal, referring to the specification by page and line number and to the drawings by reference characters, as required by 37 C.F.R.

§ 41.37(c)(1)(v). Each element of the claims is identified by a corresponding reference to the specification and drawings where applicable. Note that the citation to passages in the specification and drawings for each claim element does not imply that the limitations from the specification and drawings should be read into the corresponding claim element or that these are the sole sources in the specification supporting the claim features.

Claim 1

A method of automatic speech recognition (ASR), comprising (Figure 3 shows a method of automated speech recognition using plural different ASR engine. The method is described in connection with Figure 1: see paragraph [0031]):

receiving a speech utterance from a user (Fig. 3, #300: A participant or user (such as a telephone caller) telephones or otherwise establishes communication between communication device 40 and communication network 10. Per block 300, the communication device provides communication network 10 with an electronic input signal in a digital format. See paragraph [0031].);

assessing resources of a plurality of different ASR engines (Fig. 3, #310: Per block 310, the host computer 50 assesses the resources of the system. At this stage, for example, ports of different ASR engines 60, memory 90, database 100, or processing power of CPU 80 can be evaluated for availability. See paragraph [0032]);

assigning the speech utterance to a single ASR engine if assessing resources is within a threshold value (Fig. 3, #330: Per block 330, if the system is busy, the host computer 50, in cooperation with the resource management application 110, assigns the input signal to a single ASR engine. See paragraph [0034].);

assigning the speech utterance to a plurality of different ASR engines if assessing resources is within a threshold value (Fig. 3, # 350: On the other hand, per block 350, if the system is not busy, the host computer, in cooperation with the resource management application 110, assigns the input signal to multiple ASR engines. Here, the recognized

text from the selected ASR engines is combined to yield more accurate recognized text when compared to a single ASR engine. See paragraph [0035].); and

generating text of the speech utterance with either the single ASR engine or plurality of ASR engines (Fig. 3, #340 and 360: The assigned ASR engine or engines generate the recognized text of the input signal. See paragraphs [0034] and [0035].).

Claim 2

The method of claim 1 wherein assessing resources further comprises monitoring port utilization for each ASR engine (The host computer assesses resources of the system, and these resources include ports of different ASR engines: see paragraph [0032]).

Claim 3

The method of claim 1 wherein assessing resources further comprises evaluating processing power (The host computer assesses resources of the system, and these resources include processing power of the CPU: see paragraphs [0032] and [0045]).

Claim 4

The method of claim 1 wherein assessing resources further comprises monitoring memory utilization and input/output utilization (The host computer assesses resources of the system, and these resources include free memory and I/O usage: see paragraphs [0032] and [0045]).

Claim 8

An automatic speech recognition (ASR) system comprising (FIG. 1 shows three ASR systems 60A, 60B, and 60C. FIG. 2 shows a block diagram of an exemplary embodiment of an ASR system 60A: see paragraph [0021].):

means for processing a digital input signal from an utterance of a user (Example means is CPU 80 in Fig. 1. Per block 300 in Fig. 3, a participant or user (such as a telephone caller) telephones or otherwise establishes communication between communication device 40 and communication network 10. Per block 300, the communication device provides communication network 10 with an electronic input

signal in a digital format. See paragraph [0031].);

means for evaluating resources of the ASR system (Example means is host computer 50 in Fig. 1. The host computer 50 assesses the resources of the system. Per block 310 of Fig. 3, the host computer 50 assesses the resources of the system. At this stage, for example, ports of different ASR engines 60, memory 90, database 100, or processing power of CPU 80 can be evaluated for availability. See paragraph [0032].); and

means for selecting between a single ASR engine and a group of ASR engines to recognize the utterance of the user, wherein the means for selecting utilizes the evaluation of resources to select between the single ASR engine and the group of ASR engines (Example means is host computer 50 in Fig. 1. The host computer 50 selects between a single ASR engine or multiple different ASR engines. Per block 330 in Fig. 3, if the system is busy, the host computer 50, in cooperation with the resource management application 110, assigns the input signal to a single ASR engine. Per block 350 in Fig. 3, if the system is not busy, the host computer, in cooperation with the resource management application 110, assigns the input signal to multiple ASR engines. See paragraphs [0034] and [0035].).

Claim 9

The ASR system of claim 8 wherein the means for evaluating resources of the system monitors port utilization of the ASR engines (The host computer assesses resources of the system, and these resources include ports of different ASR engines: see paragraph [0032]).

Claim 10

The ASR system of claim 9 wherein the means for evaluating resources of the system also monitors available processing power of the system (The host computer assesses resources of the system, and these resources include processing power of the CPU: see paragraphs [0032] and [0045]).

Claim 13

The ASR system of claim 8 wherein the means for evaluating resources of the system evaluates ASR ports, system resources, and call handlers (Examination of the system resources include ASR ports, system resources, and call handlers: see paragraph [0043]).

Claim 14

A system, comprising (Fig. 1 shows a diagram of a voice telephone system, and Fig. 2 shows a diagram of an ASR system: see paragraphs [0020] and [0021]):

a computer system comprising a central processing unit coupled to a memory and resource management application (Fig. 1 shows a host computer system 50 comprising a CPU 80, memory 90, extracted algorithm 110, and resource management application 110 coupled through buses 120. See paragraph [0019]); and

a plurality of different automatic speech recognition (ASR) engines coupled to the computer system, wherein the computer system is adapted to select either a single ASR engine or multiple ASR engines to analyze a speech utterance based on resources available on the system (Fig. 1 shows plural ASR systems 60A, 60B, and 60C coupled to the computer system 50. See paragraph [0018]. The host computer 50 selects between a single ASR engine or multiple different ASR engines. Per block 330 in Fig. 3, if the system is busy, the host computer 50, in cooperation with the resource management application 110, assigns the input signal to a single ASR engine. Per block 350 in Fig. 3, if the system is not busy, the host computer, in cooperation with the resource management application 110, assigns the input signal to multiple ASR engines. See paragraphs [0034] and [0035].).

Claim 15

The system of claim 14 wherein the computer system selects an ASR engine that has most available resources (Various resources are examined to determine an engine with the most available resources: see paragraphs [0032], [0043], and [0045]).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Claims 1-5, and 8-20 are rejected under 35 USC § 102(e) as being anticipated by USPN 7,058,573 (Murveit).

Claims 6 and 7 are rejected under 35 USC § 103(a) as being unpatentable over USPN 7,058,573 (Murveit).

VII. ARGUMENT

The rejection of claims 1 – 20 is improper, and Appellants respectfully request reversal of these rejections.

The claims do not stand or fall together. Instead, Appellants present separate arguments for various claims. Each of these arguments is separately argued below and presented with separate headings and sub-heading as required by 37 C.F.R. § 41.37(c)(1)(vii).

Overview of Claims and Primary Reference (Murveit)

As a precursor to the arguments, Appellants provide an overview of the claims and the primary reference (Murveit). This overview will assist in determining the scope and content of the prior art and assist in analyzing the differences between Murveit and the claims.

The claims are directed to automatic speech recognition (ASR) methods and systems. When a speech utterance is received, the available resources of plural different ASR engines are assessed. Based on the assessed resources, the speech utterance is assigned to either a single ASR engine or plural different ASR engines if the assessment of resources is within a threshold value.

Murveit generally uses a single ASR engine to perform multiple passes for speech input. Specifically, figure 2 in Murveit shows a single speech recognition system 200 having a single processor 202. The processor performs a first speech recognition technique on the input. Then, the same processor performs a second speech recognition technique on the input (see column 2, lines 38-64). At column 7, lines 10-13, Murveit does state that the system can select a most appropriate speech recognition system from a plurality of speech recognition systems. Nowhere, however, does Murveit teach or even suggest that this selection is based on an **assessment of the resources**. Further, Murveit never discusses that a speech utterance is assigned to a plurality of different ASR engines if assessing resources is **within a threshold value**. Murveit never assesses resources in this manner.

Claim Rejections: 35 USC § 102(e)

Claims 1-5, and 8-20 are rejected under 35 USC § 102(e) as being anticipated by USPN 7,058,573 (Murveit). These rejections are traversed.

The claims recite one or more elements that are not taught in Murveit. Some examples are provided for different claim groups having identified sub-headings.

Sub-Heading: Claims 1 and 5

Claim 1 is selected for discussion of this group.

As one example, claim 1 recites assessing resources of a plurality of different ASR engines. The claim then recites assigning the speech utterance to a single ASR engine if assessing resources is within a threshold value. Murveit does not teach this element. The Examiner argues that this element is taught in Murveit at column 3, lines 4-6. Appellants respectfully disagree.

Column 3, lines 4-6 in Murveit mentions the word “threshold” but this term is used in an entirely different context than the recitations of claim 1. Specifically, Murveit teaches that a processor assigns probabilities to a spoken input after a first pass (see example in Murveit at column 2, lines 56-62 for the word “Boston”). If the probability is high that the processor accurately identified the word in the first pass, then a second subsequent pass is not necessary. In particular, if the probability is higher than a “threshold” then a second pass is not performed.

Thus, Murveit uses the word “threshold” to correlate with the probability that a processor has accurately identified a word. By contrast, claim 1 correlates the word “threshold” with an assessment of resources of ASR engines to determine whether to assign a single engine to the input speech. Again, claim 1 recites assessing resources of a plurality of different ASR engines. If the assessment of these resources is within a threshold value, then the speech utterance is assigned to the single ASR engine. Murveit never assess resources to determine whether or not to assign the input to the processor. In Murveit, the input speech appears to be automatically assigned to the processor, an assessment of resources of plural ASR engines is not first performed (see column 4, lines 26-28 stating that the input speech is provided to the ASR system).

For a prior art reference to anticipate under section 102, every element of the claimed invention must be identically shown in a single reference (see *In re Bond*, 910 F.2d 831, 15 U.S.P.Q.2d 1566 (Fed. Cir. 1990)). For at least these reasons, the claims are not anticipated by Murveit.

As another example, claim 1 recites assessing resources of a plurality of different ASR engines. The claim then recites assigning the speech utterance to a plurality of different ASR engines if assessing resources is within a threshold value. Murveit does not teach this element. The Examiner argues that this element is taught in Murveit at column 3, lines 4-6. Appellants respectfully disagree.

Column 3, lines 4-6 in Murveit mentions the word “threshold” but this term is used in an entirely different context than the recitations of claim 1. Specifically, Murveit teaches that a processor assigns probabilities to a spoken input after a first pass (see example in Murveit at column 2, lines 56-62 for the word “Boston”). If the probability is high that the processor accurately identified the word in the first pass, then a second subsequent pass is not necessary. In particular, if the probability is higher than a “threshold” then a second pass is not performed.

Thus, Murveit uses the word “threshold” to correlate with the probability that a processor has accurately identified a word. By contrast, claim 1 correlates the word “threshold” with an assessment of resources of plural different ASR engines to determine whether to assign the plural different ASR engines to the input speech. Again, claim 1 recites assessing resources of a plurality of different ASR engines. If the assessment of these resources is within a threshold value, then the speech utterance is assigned to the plural different ASR engines. Murveit never assess resources to determine whether or not to assign the input to plural different processors. In Murveit, the input speech appears to be automatically assigned to a single processor, an assessment of resources of plural ASR engines is not first performed (see column 4, lines 26-28 stating that the input speech is provided to the ASR system).

Anticipation under section 102 can be found only if a single reference shows exactly what is claimed (see *Titanium Metals Corp. v. Banner*, 778 F.2d 775, 227 U.S.P.Q. 773 (Fed. Cir. 1985)). For at least these reasons, the claims are not anticipated by Murveit.

As yet example, claim 1 recites assigning the speech utterance to a plurality of different ASR engines if assessing resources is within a threshold value. By contrast, Murveit uses a single ASR engine to perform multiple passes for speech input. Specifically, figure 2 in Murveit shows a single speech recognition system 200 having a single processor 202. The processor performs a first speech recognition technique on the input. Then, the same processor performs a second speech recognition technique on the input (see column 2, lines 38-64).

At column 7, lines 10-13, Murveit does state that the system can select a most appropriate speech recognition system from a plurality of speech recognition systems. Murveit, however, never states that these speech recognitions systems are different. In Murveit, all of the speech recognition system could be the same. Murveit is completely silent on this issue. In other words, Murveit never states that he uses plural “different” ASR engines.

Anticipation is established only when a single prior art reference discloses each and every element of a claimed invention united in the same way (see *RCA Corp. v. Applied Digital Data Systems, Inc.*, 730 F.2d 1440, 1444 (Fed. Cir. 1984)). For at least these reasons, the claims are not anticipated by Murveit.

Sub-Heading: Claims 2 and 9

Claims 2 and 9 recite that assessing resources comprises monitoring (claim 2) or evaluating (claim 9) port utilization for each ASR engine. Murveit does not teach this element. The Examiner argues that this element is taught in Murveit at column 3, lines 4-7. Appellants respectfully disagree.

Column 3, lines 4-7 in Murveit teaches that if a first pass has a high threshold (i.e., above 95%), then a second pass through the processor is not performed. This section of Murveit has nothing whatsoever to do with monitoring or evaluating port utilization of the processor. Murveit never states that his system assesses resources by evaluating or monitoring port utilization. Murveit does not even discuss port utilization.

Sub-Heading: Claims 3 and 10

Claims 3 and 10 recite that assessing resources comprises evaluating processing power. Murveit does not teach this element. The Examiner argues that this element is taught in Murveit at column 2, line 23. Appellants respectfully disagree.

Column 2, line 23 states that a second pass of the speech through the processor does not require undue processing power. By contrast, claims 3 and 10 are directed to assessing resources to determine whether to assign a speech to a single ASR engine or plural ASR engines. This assessment of resources comprises evaluating processing power.

Murveit never evaluates processing power to determine whether to assign the speech to a single or multiple ASR engines. Instead, Murveit merely makes an observation that a second pass of the speech through the processor does not require undue processing power. Murveit does not assess processing power to determine whether to send the speech to one or multiple ASR engines. Murveit mentions processing power for an entirely different reason (i.e., making a statement that a pass through the processor does not consume undue processing power).

Sub-Heading: Claim 4

Claim 4 recites that assessing the resources comprises monitoring memory utilization and input/output utilization. Murveit does not teach this element. The Examiner argues that Murveit teaches this element at column 8, lines 48-55. Appellants respectfully disagree.

Column 8, lines 48-55 in Murveit discusses that the spoken input is stored in a buffer. During the second pass, the spoken input is removed from the buffer to minimize delay. Murveit is not monitoring this buffer to assess resources for a determination whether to assign the speech to single or plural ASR engines. Instead, Murveit is using the buffer as storage for the spoken input. Murveit never assesses this buffer for determining where to send the speech utterance.

Sub-Heading: Claims 8 and 11-12

Claim 8 is selected for discussion for this group.

Claim 8 recites a means that utilizes an evaluation of resources of the ASR system to select either a single ASR engine or group of ASR engines to analyze a speech utterance. Murveit does not teach this element.

At column 7, lines 10-13, Murveit does state that the system can select a most appropriate speech recognition system from a plurality of speech recognition systems. Murveit, however, never states how this selection is made. Is the selection randomly made? Is the selection made based on the user's speech input? Is the selection made based on geography of the user? It is not possible to determine how Murveit makes this determination since the teachings in Murveit are completely silent on this aspect. In other words, Murveit never states or even suggests how an ASR engine is selected. By contrast, claim 8 expressly recites that the resources of the ASR system are evaluated to make the selection.

Anticipation is established only when a single prior art reference discloses each and every element of a claimed invention united in the same way (see *RCA Corp. v. Applied Digital Data Systems, Inc.*, 730 F.2d 1440, 1444 (Fed. Cir. 1984)). For at least these reasons, independent claim 14 and its dependent claims are allowable over Murveit.

Sub-Heading: Claim 13

Claim 13 recites that the system evaluates three different elements: ASR ports, system resources, and call handlers. Murveit does not teach this element. The Examiner argues that Murveit teaches this element at column 7, lines 17-18 and column 8, lines 48-55. Appellants respectfully disagree.

Column 7, lines 17-18 in Murveit discusses a caller's channel type. The channel type can be a caller calling from a handset. This section of Murveit has nothing whatsoever to do with evaluating the resources of the system to select between a single ASR or a group of ASR engines. Murveit never teaches that the caller type (for example,

a caller using a handset) is somehow used to determine whether the input speech from that caller will be routed to a single ASR engine or plural ASR engines.

Sub-Heading: Claims 14 and 16-20

Claim 14 is selected for discussion for this group.

Claim 14 recites that the computer system selects either a single ASR engine or multiple ASR engines to analyze a speech utterance based on resources available on the system. Murveit does not teach this element.

At column 7, lines 10-13, Murveit does state that the system can select a most appropriate speech recognition system from a plurality of speech recognition systems. Murveit, however, never states how this selection is made. Is the selection randomly made? Is the selection made based on the user's speech input? Is the selection made based on geography of the user? It is not possible to determine how Murveit makes this determination since the teachings in Murveit are completely silent on this aspect. In other words, Murveit never states or even suggests how an ASR engine is selected. By contrast, claim 14 expressly recites that the resources available on the system are evaluated to make the selection.

Anticipation is established only when a single prior art reference discloses each and every element of a claimed invention united in the same way (see *RCA Corp. v. Applied Digital Data Systems, Inc.*, 730 F.2d 1440, 1444 (Fed. Cir. 1984)). For at least these reasons, independent claim 14 and its dependent claims are allowable over Murveit.

Sub-Heading: Claim 15

Claim 15 recites that the computer system selects an ASR engine that has the most available resources. Murveit does not teach this element. The Examiner argues that Murveit teaches this element at column 3, lines 4-7. Appellants respectfully disagree.

Column 3, lines 4-7 in Murveit teaches that if a first pass has a high threshold (i.e., above 95%), then a second pass through the processor is not performed. This section of Murveit has nothing whatsoever to do with selecting an ASR engine with the most available resources. Murveit never states how an ASR engine is selected.

Claim Rejections: 35 USC § 103(a)

Claims 6 and 7 are rejected under 35 USC § 103(a) as being unpatentable over USPN 7,058,573 (Murveit). These rejections are traversed.

As noted above, Murveit does not teach or suggest all the elements of independent claim 1. Claims 6 and 7 depend from claim 1. Thus for at least the reasons provided with respect to independent claim 1, dependent claims 6 and 7 are allowable over Murveit.

CONCLUSION

In view of the above, Appellants respectfully request the Board of Appeals to reverse the Examiner's rejection of all pending claims.

Any inquiry regarding this Amendment and Response should be directed to Philip S. Lyren at Telephone No. 832-236-5529. In addition, all correspondence should continue to be directed to the following address:

Hewlett-Packard Company
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400

Respectfully submitted,

/Philip S. Lyren #40,709/

Philip S. Lyren
Reg. No. 40,709
Ph: 832-236-5529

VIII. Claims Appendix

1. A method of automatic speech recognition (ASR), comprising:
 - receiving a speech utterance from a user;
 - assessing resources of a plurality of different ASR engines;
 - assigning the speech utterance to a single ASR engine if assessing resources is within a threshold value;
 - assigning the speech utterance to a plurality of different ASR engines if assessing resources is within a threshold value; and
 - generating text of the speech utterance with either the single ASR engine or plurality of ASR engines.
2. The method of claim 1 wherein assessing resources further comprises monitoring port utilization for each ASR engine.
3. The method of claim 1 wherein assessing resources further comprises evaluating processing power.
4. The method of claim 1 wherein assessing resources further comprises monitoring memory utilization and input/output utilization.
5. The method of claim 1 wherein assessing resources further comprises monitoring a number of users providing speech utterances.

6. The method of claim 1 wherein assigning the speech utterance to a single ASR engine if assessing resources is within a threshold value occurs when port utilization of the single ASR engine is lower than a port utilization threshold of about 80%.
7. The method of claim 1 wherein assigning the speech utterance to a plurality of different ASR engines if assessing resources is within a threshold value occurs when port utilization of two ASR engines is lower than a predefined threshold of about 75%.
8. An automatic speech recognition (ASR) system comprising:
 - means for processing a digital input signal from an utterance of a user;
 - means for evaluating resources of the ASR system; and
 - means for selecting between a single ASR engine and a group of ASR engines to recognize the utterance of the user, wherein the means for selecting utilizes the evaluation of resources to select between the single ASR engine and the group of ASR engines.
9. The ASR system of claim 8 wherein the means for evaluating resources of the system monitors port utilization of the ASR engines.
10. The ASR system of claim 9 wherein the means for evaluating resources of the system also monitors available processing power of the system.

11. The ASR system of claim 8 further comprising a means for combining results of ASR engines if the group of ASR engines is selected, the group of ASR engines being adapted to provide a more accurate recognition of the utterance than a single ASR engine.

12. The ASR system of claim 8 wherein the means for evaluating resources of the system evaluates resources to simultaneously run multiple ASR engines.

13. The ASR system of claim 8 wherein the means for evaluating resources of the system evaluates ASR ports, system resources, and call handlers.

14. A system, comprising:

a computer system comprising a central processing unit coupled to a memory and resource management application; and

a plurality of different automatic speech recognition (ASR) engines coupled to the computer system, wherein the computer system is adapted to select either a single ASR engine or multiple ASR engines to analyze a speech utterance based on resources available on the system.

15. The system of claim 14 wherein the computer system selects an ASR engine that has most available resources.

16. The system of claim 14 further comprising a telephone network comprising at least one switching service point coupled to the computer system.

17. The system of claim 16 further comprising at least one communication device in communication with the switching service point to provide the speech utterance.

18. The system of claim 14 wherein the resource management application comprises a recognition proxy component and a resource monitoring component.

19. The system of claim 18 wherein the resource management component collects and analyzes information about the resources available on the system.

20. The system of claim 19 wherein the resource monitoring component mediates between the plurality of ASR engines and the resource management component.

IX. EVIDENCE APPENDIX

None.

X. RELATED PROCEEDINGS APPENDIX

None.